

Control the flow of ideas online and you control the truth, finds Chris Baraniuk

WORLD WIDE WARP

OPERATION Jade Helm 15 wouldn't start for another few weeks. But the footage of armed troops advancing down suburban streets confirmed many people's fears. This was evidence of the Obama administration's hushed-up plan to wage war on the American people. "They're trying to incite violence, the government, so that they can enact martial law in this country," one YouTuber explained. Or was it a plot to confiscate firearms? Or a plan to invade Texas?

Or none of those things. When or where the video was shot was unclear – the date stamp read "2073". But if you were inclined to believe the conspiracy theories about Jade Helm 15, a routine large military training exercise, you would have found no shortage of material online. A few years ago, it might only have bounced around remote corners of the internet. Instead, the rumours spread via social media so quickly and widely that Texas governor Greg Abbott ordered the state guard to monitor the military's activities.

From Islamic State's recruitment drives to the skilful use of social media by companies and individuals to promote their brands and ideas, information and misinformation are increasingly hard to distinguish online – and none of us are immune. "There's something happening here that's really unprecedented," says Robert Epstein, a psychologist at the American Institute for Behavioral Research and Technology in Vista, California. "Technologies are rapidly evolving that can impact people's behaviours, opinions, attitudes, beliefs on a

massive scale – without their awareness." How can we help the truth to hit back?

Mass delusions are not new, of course. The first radio broadcast of H. G. Wells's *The War of the Worlds* in 1938 triggered widespread panic. Thousands of people jammed emergency lines believing that Martians were invading. It is easy to scoff. But in many ways we are just as credulous today. Quirks of the social web can make falsehoods spread far more widely and more quickly. What's more, unlike the one-off worry of an alien invasion, this misinformation can change long-term opinions. Indeed, some worry that the internet is turning into the biggest mind control experiment the world has seen. You may think you're savvy, but there are armies of people out there equipped with technology, all promoting their own version of reality.

Marketers, lobbyists, activists, extremists – they all depend on being able to sway opinion. And with the social web it's easier than ever. The World Economic Forum ranks massive digital misinformation as a geopolitical risk alongside terrorism and failure of global governance.

Politicians, too, are learning how to use online promotion tools to their advantage. The popularity of controversial US presidential candidate Donald Trump, for example, is in part due to social media campaigns. These examples have real world impacts, influencing democratic processes or financial markets. "We're entering an era of unprecedented psychological manipulation," says Bruce

Schneier, a director of the Electronic Frontier Foundation in San Francisco. But the manipulation can be subtle and often hard to notice at all (see "Warped web", page 40).

To get a better understanding of what manipulation is possible, researchers have started to study how the social web works: what's said, how it's said, who says it. Understand that and you may be able to predict how information spreads online and control the flow of ideas.

Walter Quattrocchi at the IMT Institute for Advanced Studies in Lucca, Italy, and his colleagues looked at how different types of information are spread on Facebook by different communities. Specifically, the researchers analysed two groups: those who shared conspiracy theories and those who shared science news articles. They found that

"We're entering an era of unprecedented psychological manipulation"

science stories received an initial spike of interest and were shared or "liked" frequently. Conspiracy theories started off with a low level of interest but sometimes grew to be even more popular than the science stories overall.

More importantly, both groups tended to ignore information that challenged their views. In another study of 55 million Facebook users, the team found that out of 50,000 posts debunking rumours, only




WARPED WEB: 5 WAYS THE NET TWISTS THE TRUTH

MAJORITY ILLUSION


Groups within social networks Facebook, Twitter and Digg tend to be influenced most by a small number of highly popular individuals, says Kristina Lerman at the University of Southern California. She studied the links between group members of social networks, and found that the actions of these well connected individuals – perhaps when sharing a politically charged post or adding an activist symbol to their profile picture – were often interpreted by the others as indicative of a consensus. Lerman has dubbed this “the majority illusion” (see diagram, right). People are generally not aware of how well connected the people they follow on social media are. In turn this makes it difficult to evaluate how popular certain activities or opinions may be.

This is a problem when the opinions being voiced by a popular individual are dangerous or incite hatred. But it may also have contributed to the rapid shift in popular opinion during the Arab Spring or concerning gay marriage in the West. “These popular people can change behaviours,” says Lerman.

The structure of social networks can make minority views seem more popular than they are



A few people sharing widely gives the impression that everyone is thinking along the same lines



If the sharers have fewer connections, the information seems much less popular

LIST ORDERING

If you want to get elected, make sure you’re good at search engine optimisation. Robert Epstein, a psychologist at the American Institute for Behavioral Research and Technology in Vista, California, and his colleague Ronald Robertson have found that the higher up the list of search results you come, the more likely people are to think you’re a credible choice. In one experiment, they found the number of people who said they intended to vote for a candidate increased by more than 48 per cent when using a biased search engine that ranked the candidate higher. The effect was lower when tested during an actual election in India, but still potentially large enough to swing a close poll.

We’ve long known that coming top of a list can have an influence. But this is different. “There is no other list phenomenon like this,” says Epstein. Because we constantly use search engines like Google to check facts – such as the time in Sydney, or currency exchange rates – we are conditioning ourselves to trust the top results. “Over and over again in routine searches we are learning that what is at the top is best,” says Epstein.

about 1 in 12 reached people who had shared the rumour in question. It is a good example of confirmation bias leading to an echo chamber, says Quattrociocchi. Information that does not fit with an individual’s world view does not get passed on. And on social networks, people trust their peers and use them as their primary information sources. “The role of the expert is going to disappear,” he says.

Governments are also trying to get to grips with the spread of information online. DARPA, the US military’s research agency, has poured money into its Social Media in Strategic Communication programme, which funded dozens of studies looking at everything from

subtle linguistic cues in specific posts to how information flows across large networks.

For example, Eric Gilbert and Tanushree Mitra at the Georgia Institute of Technology in Atlanta analysed 45,000 projects on crowdfunding site Kickstarter to see what made some succeed and others fail. They found that language alone accounted for more than half of the variance between successful and unsuccessful projects.

At a network level, it has previously been suggested that information spreads much like a contagious disease – in that having more contacts makes exposure more likely. However, Nathan Hodas and Kristina Lerman


at the University of Southern California Information Sciences Institute in Marina del Rey found the opposite. Studying the flow of ideas on Twitter and Digg – another social network – revealed that highly connected people are less likely to see a given piece of information. Once infected, however, they have a greater impact on their followers.

Their model allows real-time forecasting of the parts average users play in a specific social network as information spreads across it. Another DARPA project showed automatic analysis of Twitter users’ activity can identify those most likely to retweet information on a given topic. These users can then be targeted and asked to share specific information.

Digital propaganda

Ultimately, the aim of such research is to find ways to identify misinformation and effectively counter it, reducing the ability of groups like ISIS to manipulate events. “They have managed to digitise propaganda in a way that is completely understanding of social media and how it’s used,” says Jonathan Russell, head of policy at counter-extremism think tank Quilliam in London.

A lack of other voices also gives the impression that they are winning, says Russell. “There’s no other effective media coming out



Ideas are more seductive when shared by friends

ASTROTURFING

One way to get your point across is to invent a wave of support or dissent. Advertisers, political groups and even governments have been accused of this.

Last year, for example, Chinese journalist Chai Jing released a documentary called *Under the Dome*, which suggested that pollution in China's cities was the reason her unborn daughter had developed a benign tumour. The video went viral, along with many comments endorsing her view. It wasn't long, though, before the Chinese government had the documentary removed from various websites, and the move coincided with a wave of negative social media posts about it. David Holmes at Monash University in Melbourne, Australia, says it's likely these were written by the government-sponsored "50 cent party", a group believed to be paid for their posts.

"The ability to understand who your audience is precisely and target them is unprecedented and dangerous," says Bruce Schneier of the Electronic Frontier Foundation. He thinks this kind of manipulation should be made illegal.

BOT DEMOCRACY

Social media accounts are sometimes controlled by software. An army of online bots can be programmed to give the impression of grassroots support for a policy, for example.

There have also been reports of bots being used by candidates in US presidential elections and a parliamentary candidate in London admitted to using the tactic in 2012.

Philip Howard at the University of Washington is about to embark on a five-year project with the Oxford Internet Institute to study how bots might influence European elections. Howard deploys bots on social networks himself to study their behaviour. Sometimes they even disagree with each other. On one occasion, one of Howard's dummy accounts argued the case for childhood vaccination with one that was against it.

"Algorithms and bots are bad for democracy if the scripts behind them are hidden, and this is the problem we have now," Howard says. "At least in Las Vegas the code behind gambling machines is audited by the state."

DISINFORMATION

Deliberately planted false information can have drastic effects. In July 2014, Cynk Technology briefly became worth \$6 billion, after a barrage of tweets and emails promoting the stock flooded the web. It turned out that the firm had no assets, no revenue and one employee.

And in November, the US Securities and Exchange Commission indicted a 62-year-old Scottish man for bogus tweets that caused stock in two other tech firms to crash, allowing him to buy their shares at knock-down prices. Such hype need not even reach real humans, says Emilio Ferrara at the University of Southern California. "You can trick automatic trading bots to buy or sell particular stocks."

Governments can also play the disinformation game. Classified documents from 2011, recently published by website *The Intercept*, appear to detail tactics used by the UK's GCHQ intelligence agency to target online fraud and groups it considers to have extremist views. These include "setting up Facebook groups, forums, blogs and Twitter accounts that encourage and monitor discussion on a topic".



MIKHAIL POCHUEV/GETTY IMAGES/PHOTO SHOOT

After a passenger jet was shot down in 2014, Russia and Ukraine spread their own versions of events online

Google and Facebook are developing "fact checker" algorithms that will continuously monitor information being shared on their networks and flag things deemed untrue.

But even if such techniques prove effective – or perhaps especially if they do – it raises difficult questions. Rand Waltzman, the ex-director of DARPA's social media research program, has argued that the US government should rethink its policy of not using social networks to influence public opinion. But do we want governments – or tech firms like Google – to become the arbiters of truth? And if not, to whom should that role belong?

These are questions we will have to face up to sooner rather than later. A recent survey by Ofcom, the UK's telecommunications regulator, showed that young people were more likely to trust information from news and social websites in 2015 than they were the previous year, for example.

Metaxas believes we have entered an era in which the definition of literacy needs to be updated. "In the past to be literate you needed to know reading and writing," he says. "Today, these two are not enough." Information reaches us from a vast number of sources. We need to learn what to read, as well as how. ■

of Iraq and Syria." Quilliam has attempted to counter such narratives with videos like *Not Another Brother*, which depicts a jihadist recruit in desperate circumstances. It aims to show how easily people can be seduced by exposure to a narrow view of the world.

Of course, governments are out to control the flow of ideas, too. Europe and the US are putting pressure on tech giants to police their communities. At a meeting of technology firms and government officials in London in January, Google recommended shutting down social media accounts linked to extremists as soon as they are created. But the danger is that those accounts will simply

move to unpoliced networks.

Misleading or dangerous posts can at least be flagged up by users. Most social networks rely on some degree of self-censoring already and this can also be automated. Panagiotis Metaxas of Wellesley College in Massachusetts has helped build a tool called Twitter Trails, which analyses the trustworthiness of stories being discussed on the network. It works by looking for telltale signs that a breaking story is false, such as it not being retweeted widely and tweets expressing scepticism.

DARPA has also sponsored a challenge to design bots that can sniff out misinformation deliberately planted on Twitter. And both

Chris Baraniuk is a science and technology journalist based in London